

REMARKS/ARGUMENTS

Claims 1-34 are pending in the application.

Claims 5-7, 33 and 34 are rejected under 35 USC §112, second paragraph, as being indefinite for failing to particularly point out and distinctly claim the subject matter which applicant regards as the invention. It is noted with appreciation that claims 1-34 are otherwise allowable over the prior art.

A requirement for information pursuant to 37 C.F.R. §1.105 was made by the examiner.

Claims 35-50 have been withdrawn by the examiner. Pursuant to MPEP §708.07 VIII, Applicant affirms withdrawal of claims 35-50 without traverse.

As to the Section 112 rejection of claims 5 and 7, a typographical error resulted in incorrect dependency in each claim. Claims 5 and 7, each depend from claim 1, and the claims have been amended accordingly.

As to the Section 112 rejection of claim 33, the claim recites “a file stored in the storage system.”

As to the Section 112 rejection of claim 34, the extraneous “wherein” term has been canceled.

Discuss of Additional Information Pursuant to 37 C.F.R. §1.105

A request was made to for a discussion of which claim features, if any, recited in independent claims 1, 8, 18, 28, and 33 correspond to prior art elements disclosed in the background section of the specification, including the mentioned PCT publication WO 99/38093. Specifically, the examiner requested a discussion of the claim limitations recited among the independent claims and summarized by the examiner as follows: (1) data object divisible into one or more partitions; (2) producing one or more replicas of a partition if there are none; and (3) producing one or more replicas of a partition if the number of existing replicas is less than a threshold.

A. Background Discussion in the Specification

1. RAID storage systems

Page 1, lines 27-28 in the specification describe RAID storage systems. “RAID” originally stood for Redundant Array of Inexpensive Disks; though some contemporaries use the term “independent” in place of “inexpensive”. Different RAID levels are defined. Some RAID levels store data (e.g., a file) across multiple disks, called “data striping” in RAID nomenclature, where portions of the file are stored in different disks. However, RAID level definitions do not involve “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

2. PCT publication No. WO 99/38093

Page 2, line 8 in the specification identifies PCT publication No. WO 99/38093. The reference observes that data information must often be identified to be in a particular state, denoted by the state of an asset, such as data files, multimedia files and fragments, records from structured databases, etc. *Page 1, lines 16-21*. A problematic attribute of digital information is that copies may exist which are identical in content but differ in the meta data that the computer uses to relate to the digital information. *Page 2, lines 3-7*.

The reference discloses a method for representing digital information in an electronic paper clip (e-CLIP) This is a reproducible identifier for a collection of digital information. The e-CLIP may represent a file, a group of files, or collections data such as selected database records, selected frames from digital or audio streams, messages from streams, etc. *Page 7, lines 9-19*. The e-CLIP can be useful for identifying groups of files that have been backed up. *Page 9, lines 17-19*.

The reference discloses a cryptographic hash function to compute an identifier for the data being represented. *Pages 12-17*. The reference discusses how an e-CLIP is generated (pages 17-21) and how an e-CLIP is retrieved (pages 21-28). There is no discussion relating to producing replicas of partitions of data objects.

B. Discussion of IDS References

1. U.S. Pat. No. 5,440,727 to Bhide

Bhide discloses a partitioned database in which one or more secondary replicas of each partition is maintained by asynchronously sending modified pages from the primary replica (P) to the secondary replica (S). *Abstract.* Bhide describes the notion of failure-free operation beginning in column 8, line 23. Database operations are executed only on the primary replica. In order to keep replicas P and S reasonably synchronized with respect to the database state, updated pages are sent to S at some time before they are discarded from P's buffer. Bhide notes there are a number of policies possible for doing this. They involve different trade-offs between recovery time and CPU, disk and network overheads during failure-free processing depending on how soon after update (i.e., dirty) pages are sent to S. Bhide discloses the following procedure to update the secondary replica (column 8, lines 50-67):

```
For each dirty page in buffer
  if (S is up)
  {
    Send page to S after latest log record modifying
    the page is written out and before the page is
    expelled from the buffer.
    After receiving acknowledgement from S that page is on
    disk delete entry for page from DPT for S.
  }
  else
    /* Do nothing, i.e. P behaves as if it is the only replica */
}
At checkpoint time do:
  if (S is up)
  {
    write DPT of both P and S in checkpoint.
    inform S that a checkpoint has taken place.
    /* For S to reset the "Received List" described later
  }
  else
    write DPT of P and a pointer to the latest SDownCP.
```

Though Bhide discloses “partitioned databases”, there is no discussion of “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

2. U.S. Pat. No. 5,778,395 to Whiting

Whiting discloses a backup process in which duplicate files (or portions of files) may be identified across nodes of a computer network, so that only a singly copy of the contents of the duplicate files (or portions thereof) is stored in the backup storage means. *Abstract*. The backup process is disclosed beginning in column 7, line 6. The process includes backing up directory files and backing up data files. *Col. 8, lines 42-44*. The contents of the backup directory file indicate the directory structure of the source disk volume. *Id at lines 46-48*. The backup data file contains data from the files included in a backup set. *Col. 13, lines 10-25*. Some of this data may be represented by references into other backup data files from previous backups, either from this user or another user. If user A has an exact copy of a file that has already been backed up by user B, user A's <fileEntry> 207 will contain the identical <fileID> 214 as user B's, but they will have distinct <dirItemNum> values 223, which are not shared between users.

Whiting, however, does not show "producing one or more replicas of a partition if there are none" or "producing one or more replicas of a partition if the number of existing replicas is less than a threshold." Quite the contrary, Whiting teaches avoiding duplicated files.

3. U.S. Pat. No. 5,815,649 to Utter

Utter shows a fault-tolerant data storage system in which storage nodes store data in at least one replicated partition group. *Abstract*. Operation of a storage node is described beginning in column 6, line 62. The discussion given in column 6, line 67 to column 7, line 34 describes partitions and storing partitions. Utter, however, does not describe "producing one or more replicas of a partition if there are none" or "producing one or more replicas of a partition if the number of existing replicas is less than a threshold."

4. U.S. Pat. No. 6,192,472 to Garay

This reference discloses partitioning information in the context of data reconstruction for fault recovery in information dispersal methods. The reference, however, does not discuss "producing one or more replicas of a partition if there are none" or "producing one or more replicas of a partition if the number of existing replicas is less than a threshold."

5. U.S. Pat. No. 6,405,315 to Burns

Burns discloses a decentralized file system on a network of remotely encrypted storage devices. Files or directories are composed of one or more streams, which logically partitions the data associated with the files or directories. *Abstract*. Burns discloses that network object are used to store the data. *Col. 6, lines 21-32*. File data objects contain all or part of the file data of a file. Directory data objects contain all or part of the directory entries of a directory. More specifically, files and directories are made up of one or more streams. Each stream logically partitions the associated data, and is composed of the data objects. *Id at lines 62-65*. Burns, however, does not show “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

6. U.S. Publ. No. 2002/0032691 to Rabii

Rabii describes a non-hierarchical directory structure for a mass storage unit. The disk is partitioned into segments of equal size. Data objects reside wholly and contiguously within a given area of the disk segments. A given object is not allowed to occupy more than one segment. During a storage operation, objects are assigned to segments in a round-robin fashion, to equalize segment utilization. *Paragraph [0012]*. Fig. 2 shows objects are stored object data partitions, where each data object resides wholly within one segment, and no object is allowed to span segments. *Paragraphs [0038-0039]*.

Paragraphs [0084-0086] discuss storing new objects. There is no discussion of “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

7. U.S. Publ. No. 2002/0194209 to Bolosky

Bolosky shows a file format for a file system in which the data stream of the file is divided into multiple blocks. *Paragraph [0007]*. As described in paragraph [0037], the files stored by the file system are distributed among the various devices 102-106 (Fig. 1) and stored in encrypted form. When a new file is created, the device on which the file is being created

encrypts the file prior to communicating the file to other device(s) for storage. The directory entry (e.g., the file name) for a new file is also communicated to the other device(s) for storage. Additionally, if a new folder or directory is created, the directory entry (e.g., folder name or directory name) is also communicated to the other device(s) for storage.

Paragraph [0145] describes the construction of a file in the file system, and paragraphs [0161-0163] describe writing to a file. Notably, there is no discussion of “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

8. U.S. Publ. No. 2002/0200207 to Dickinson

Dickinson discloses an update file for transmission to a client computer that permits the client computer to generate a copy of a current version of a subscription file from a copy of an earlier version of the subscription file. For each segment of the current version of the subscription file, the server computer searches an earlier version of a signature list for an old segment signature which matches a new segment signature corresponding to the segment. When a match is detected, the server computer writes a command in the update file for the client computer to copy an old segment of the client computer's copy of the earlier version of the subscription file into the client computer's copy of the current version of the subscription file, where the old segment corresponds to the segment for which a match was detected. When no match is detected, the server computer writes a command into the update file for the client computer to insert a new segment of the current version of the subscription file into the client computer's copy of the current version of the subscription file, where the new segment of the current version of the subscription file is written into the update file. The update file is only generated when the server computer determines the subscription file has changed. The server computer periodically monitors the subscription file to determine if it has been altered before generating an update file. *Abstract and paragraphs [0021-0023]*.

In Fig. 9 and paragraph [0061] Dickinson describes the creation of a current copy on the client computer of the current subscription file from a copy of the earlier version of the subscription file on the client computer using the update file 801 shown in FIG. 8. Dickinson,

Appl. No. 10/801,898
Amdt. sent March 24, 2006
Reply to Office Action of December 19, 2005

PATENT


however, does not show “producing one or more replicas of a partition if there are none” or “producing one or more replicas of a partition if the number of existing replicas is less than a threshold.”

CONCLUSION

In view of the foregoing, Applicants believe all claims now pending in this Application are in condition for allowance. The issuance of a formal Notice of Allowance at an early date is respectfully requested.

If the Examiner believes a telephone conference would expedite prosecution of this application, please telephone the undersigned at 650-326-2400.

Respectfully submitted,


George B. F. Yee
Reg. No. 37,478

TOWNSEND and TOWNSEND and CREW LLP
Two Embarcadero Center, Eighth Floor
San Francisco, California 94111-3834
Tel: 650-326-2400
Fax: 415-576-0300
GBFY:gjs
60674627 v1